

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPEAL BRIEF

Applicant:	Birkestrand, <i>et al.</i>	Docket No.:	ROC920030150US1
Serial No.:	10/616,675	Group Art Unit:	2195
Filed:	07/10/03	Examiner:	Zhe, Meng Yao
TITLE:	ASSURING RECOVERY OF TEMPORARY RESOURCES IN A LOGICALLY PARTITIONED COMPUTER SYSTEM		

Mail Stop APPEAL BRIEF - PATENTS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir/Madam:

This appeal is taken from the Examiner's final rejection, set forth in the Office Action dated 09/04/08, of appellant's claims 10-11. Appellant's Notice of Appeal under 37 C.F.R. § 1.191 was filed on 12/04/2008.

REAL PARTY IN INTEREST

International Business Machines Corporation is the Real Party in Interest.

RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences for this patent application.

STATUS OF CLAIMS

Claims 1-24 were originally filed in the patent application. In response to the first office action dated 02/28/2007, an amendment was filed on 05/29/2007 that cancelled claims 3, 9, 15-16, 18, and 20-21 and amended claims 1, 4, 8, 10, 14, and 19. In response to the final office action dated 08/24/2007, an RCE and amendment were filed on 11/26/2007 that cancelled claims 1-2, 6-8, 12-14, 17, and 23-24 and amended claims 4, 10 and 19. In response to a third office action dated 02/20/2008, an amendment was filed on 05/20/2008 that cancelled claims 4-5, 19, and 22 and amended claims 10-11. In the pending final office action dated 09/04/2008, claims 10-11 were rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 7,136,800 to Vega in view of U.S. Patent No. 7,146,492 to Circenis *et al.* (hereinafter "Circenis"). Claims 10-11 were also rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 7,032,222 to Karp *et al.* (hereinafter "Karp") in view of Vega and Circenis. No claim was allowed. The claims at issue in this appeal are claims 10-11, all of which stand finally rejected.

STATUS OF AMENDMENTS

After the first office action dated 02/28/2007, an amendment was filed on 05/29/2007, which has been entered. After the second office action, an RCE and amendment were filed on 11/26/2007, which has been entered. After the third office action, an amendment was filed on 05/20/2008, which has been entered. Therefore, the claims at issue in this appeal are claims 10-11 as amended in the amendment filed 05/20/2008.

SUMMARY OF CLAIMED SUBJECT MATTER

Claim 10 recites a computer-implemented method for providing at least one temporary resource on demand for a specified resource-time in a computer system that includes a plurality of logical partitions, the method comprising the steps of: requesting an enablement code corresponding to the at least one temporary resource for the specified resource-time (step 410 in FIG. 5; p. 12 lines 18-19); receiving the enablement code (step 420 in FIG. 5; p. 12 lines 19-20); enabling the at least one temporary resource for the specified resource-time (step 430 in FIG. 5; p. 12 lines 22-23); using the at least one temporary resource for the specified resource-time (steps 440, 450 and 460 in FIG. 6; p. 12 lines 23-26); reading a minimum resource specification for all of the plurality of logical partitions, each minimum resource specification specifying minimum resources that must be available for a corresponding logical partition to function correctly (210 in FIG. 2; p. 11 lines 1-2); one of the plurality of logical partitions requesting to increase a minimum resource specification corresponding to the one logical partition (FIGS. 6 & 7; p. 15 lines 10-13); determining a total of resources that are permanently enabled in the computer system (FIGS. 6 & 7; p. 15 lines 15-16; p. 8 lines 9-10); determining whether the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all of the plurality of logical partitions to exceed the total of resources that are permanently enabled in the computer system (FIGS. 6 & 7; p. 15 lines 10-13); if the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all logical partitions to exceed the total of resources that are permanently enabled in the computer system, denying the requested increase in the minimum resource specification corresponding to the one logical partition (FIGS. 6 & 7; p. 15 lines 8-16); and if the requested increase in the minimum resource specification corresponding to the one logical partition will not cause the sum of all minimum resource specifications for all logical partitions to exceed the total of resources that are permanently enabled in the computer system, performing the steps of: permitting

the requested increase in the minimum resource specification corresponding to the one logical partition (FIGS. 6 & 7; p. 15 lines 13-14); and enabling the at least one temporary resource (step 430 in FIG. 5; p. 12 lines 22-23); recovering the at least one temporary resource when the specified resource-time expires (step 470 in FIG. 5; p. 12 lines 25-26).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The following grounds of rejection are presented for review on this Appeal:

- 1. Whether claims 10-11 are unpatentable under 35 U.S.C. §103(a) as being unpatentable over Vega in view of Circenis.**
- 2. Whether claims 10-11 are unpatentable under 35 U.S.C. §103(a) as being unpatentable over Karp in view of Vega and further in view of Circenis.**

ARGUMENT

Issue 1: Whether claims 10-11 are unpatentable under 35 U.S.C. §103(a) as being unpatentable over Vega in view of Circenis.

Claim 10

In rejecting claim 10, the examiner reads col. 6 lines 1-12 of Vega on the minimum resource specification for each of the plurality of logical partitions. Note, however, the guaranteed minimum capacity fraction in Vega is only assigned to the most critical of the virtual machines of the computer system. See Vega at co. 6 lines 39-43. This means some of the virtual machines are not assigned a guaranteed minimum capacity fraction. For this reason, the guaranteed minimum capacity fraction taught in Vega cannot read on the minimum resource specification in claim 10 that is defined for all of the logical partitions and that specifies minimum resources that must be available for a corresponding logical partition to function correctly.

In addition, claim 10 recites:

. . . one of the plurality of logical partitions requesting to increase a minimum resource specification corresponding to the one logical partition;
. . .

In Vega, the guaranteed minimum capacity fractions must be entered manually. See Vega at col. 6 lines 26-27. There is no teaching in Vega of a logical partition requesting to increase its minimum resource specification. The examiner admits Vega does not teach the limitation recited above and states:

However, Vega does teach contending for resources in general among logical partitions for the purpose of gaining more resources (7, line 60 – Column 8, line 6). It would have been obvious to one having ordinary skill in the art to modify the teachings of Vega such that minimum

resources specifically, not just resources in general, may also be contended among logical partitions such that the logical partition may request for an increase in minimum resources as well, because it allows for the gain of more resources.

The examiner has confused two very separate and distinct concepts in Vega, and has assumed they are somehow related. The contention of resources referenced by the examiner at col. 7 line 60 to col. 8 line 6 in Vega relates to run-time contending for resources among virtual machines. This is abundantly clear from co. 7 lines 57-61, which states:

Once the allocation of processor resources has been determined by one of the techniques described herein, meta-scheduler 15 modulates each virtual machine's access to processor resources according to the prescribed utilization goals and minimum access requirements described above.

The lines that follow that were cited by the examiner relate to the run-time modulation of the virtual machine's access to processor resources. Thus, Vega discloses two distinct phases: 1) allocation of resources prior to run-time; and 2) modulation of resources at run-time based on allocation in phase 1. These two phases occur at different times, so applying the teachings of phase 2 to phase 1 as suggested by the examiner would not have been obvious to one of ordinary skill in the art at the time the invention was made.

The statement by the examiner quoted above as allegedly justifying extending Vega to a logical partition requesting to increase a minimum resource specification corresponding to the one logical partition is riddled with problems. First of all, the examiner appears to have confused the claim term "minimum resource specification" with "minimum resource specifically." These are clearly not the same. Second, the examiner has taken the teachings of Vega regarding run-time modulation of resources to virtual machines (in phase 2 described above) and stated it would have been obvious for logical partitions to contend for these resources during allocation of resources (in phase 1

described above). The run-time contending of processor resources in phase 2 of Vega has nothing whatsoever to do with the capacity fractions for the virtual machines that are manually entered during phase 1. As a result, the examiner's contention that it would be have been obvious to one of ordinary skill in the art at the time the invention was made to apply the run-time modulation of resources in phase 2 to the allocation of resources in phase 1 to arrive at "one of the plurality of logical partitions requesting to increase a minimum resource specification corresponding to the one logical partition" is clear error. Vega teaches entering capacity fractions for each virtual machine manually. The run-time modulation of processor resources during phase 2 of Vega cannot affect the manually-entered capacity fractions during phase 1. In fact, the run-time modulation in phase 2 is done consistent with the manually-entered capacity fractions in phase 1. Neither the manually-entered capacity fractions in phase 1 nor the run-time modulation in phase 2 nor their combination teach or suggest one of the plurality of logical partitions requesting to increase a minimum resource specification corresponding to the one logical partition as recited in claim 10. For this reason alone, claim 10 is allowable over the combination of Vega and Circenis.

Claim 10 also recites:

- ... determining a total of resources that are permanently enabled in the computer system;

- determining whether the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all of the plurality of logical partitions to exceed the total of resources that are permanently enabled in the computer system;

- if the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all logical partitions to exceed the total of resources that are permanently enabled in the computer system, denying the requested increase in the minimum resource specification corresponding to the one logical partition; and

- if the requested increase in the minimum resource specification corresponding to the one logical partition will not cause the sum of all

minimum resource specifications for all logical partitions to exceed the total of resources that are permanently enabled in the computer system, performing the steps of:

 permitting the requested increase in the minimum resource specification corresponding to the one logical partition; and
 enabling the at least one temporary resource; . . .

The examiner cites to column 6 lines 18-33 of Vega as allegedly teaching all of the above limitations. Vega does teach a reserve or minimum amount of processor resources (referred to in Vega as a “capacity fraction”) for a virtual machine and the resources above that reserve amount are shared with other virtual machines, but nowhere does Vega teach or suggest increasing the capacity fraction and especially does not teach or suggest a logical partition or virtual machine requesting to increase the capacity fraction. Vega expressly states the capacity fractions must be entered manually at col. 6 lines 26-27, which goes entirely against the examiner’s assumption that a logical partition could request to increase the minimum resource specification for the logical partition. In addition, nowhere does Vega distinguish between resources that are enabled permanently or that are enabled temporarily. For the many reasons given above, claim 10 is allowable over the combination of Vega and Circenis, and appellant respectfully requests the examiner’s rejection of claim 10 under 35 U.S.C. §103(a) based on the combination of Vega and Circenis be reversed.

Claim 11

Claim 11 depends on claim 10 which is allowable for the reasons given above. As a result, claim 11 is allowable as depending on an allowable independent claim. Appellant respectfully requests the examiner’s rejection of claim 11 under 35 U.S.C. §103(a) based on the combination of Vega and Circenis be reversed.

Issue 2: Whether claims 10-11 are unpatentable under 35 U.S.C. §103(a) as being unpatentable over Karp in view of Vega and further in view of Circenis.

Claim 10

Claim 10 recites:

... reading a minimum resource specification for all of the plurality of logical partitions, each minimum resource specification specifying minimum resources that must be available for a corresponding logical partition to function correctly; ...

In rejecting this limitation, the examiner cites col. 3 line 57 to col. 4 line 8 in Karp. This portion of Karp relates to a resource manager 12 determining whether to allocate $n1$ units requested by a task. It is not clear from the rejection which teaching in Karp allegedly reads on the minimum resource specification in claim 10. Perhaps the examiner believes the soft limit or hard limit in Karp reads on the minimum resource specification. The soft limit and hard limit in Karp relate to users, not tasks. This is clear from the first three lines of the Karp Abstract, which states:

A method for flexible allocation of a resource in which a soft limit and a hard limit are assigned to each of a set of potential users of the resource.

This means the limits in Karp apply to users, not tasks as assumed by the examiner.

Specifying limits on resources for users does not read on a minimum resource specification for a logical partition. In addition, the minimum resource specification in claim 10 specifies minimum resources that must be available for a corresponding logical partition to function correctly. The soft limit and hard limit in Karp apply to a user, and do not specify minimum resources. As such, neither the soft limit nor the hard limit in Karp specify “minimum resources that must be available for a corresponding logical

partition to function correctly” as recited in claim 10. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

Claim 10 further recites:

... determining a total of resources that are permanently enabled in the computer system; ...

The examiner did not specifically map any teaching of Karp on this limitation, but instead lumped it with the next limitation, and cited col. 3 lines 44-51 of Karp as allegedly reading both. The cited portion of Karp relates to a high watermark and a low watermark. The high watermark is an upper limit on the total utilization of the resource 10. Karp at col. 3 line 45-46. Nowhere does this portion or any other portion of Karp teach or suggest determining a total of resources that are permanently enabled in the computer system. In fact, a text search of Karp show no occurrence of the words “permanent” or “permanently”. Nowhere does Karp teach or suggest determining a total of resources that are permanently enabled in the computer system. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

Claim 10 further recites:

... determining whether the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all of the plurality of logical partitions to exceed the total of resources that are permanently enabled in the computer system; ...

In rejecting this limitation, the examiner cites to col. 3 lines 44-51 of Karp. The cited portion of Karp relates to a high watermark and a low watermark. The high watermark is an upper limit on the total utilization of the resource 10. Karp at col. 3 line 45-46. The high watermark in Karp does not relate to resources that are permanently enabled in the computer system. The cited language in Karp has nothing that can be reasonably read on “a sum of all minimum resource specifications for all of the plurality of logical partitions

to exceed the total of resources that are permanently enabled in the computer system” as recited in claim 10. In Karp, where are the minimum resource specifications for all of the plurality of logical partitions? Where are the total of resources that are permanently enabled in the computer system? Simply put, they don’t exist in Karp. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

Claim 10 further recites:

... if the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all logical partitions to exceed the total of resources that are permanently enabled in the computer system, denying the requested increase in the minimum resource specification corresponding to the one logical partition; ...

In the rejection, the examiner states Karp teaches these limitations, citing col. 4 lines 30-36 of Karp. The cited portion of Karp teaches things that may, at first blush, appear somewhat similar to the limitations in claim 10 quoted above. For the examiner’s rejection of this clause in claim 10 to have merit, the request for n1 units in Karp would have to read on the requested increase in the minimum resource specification corresponding to the one logical partition; the total allocation of resources in Karp would have to read on the sum of all minimum resource specifications for all logical partitions; and the high watermark in Karp would have to read on the total resources that are permanently enabled in the computer system. In fact, the teachings in Karp read on none of these three specific limitations enumerated above for the reasons given in detail below.

The request for n1 units in Karp is a request by a task to allocate an additional N1 units to a user that is running the task. The request for the additional n1 units in Karp does not affect the soft limit or the hard limit assigned to the user. The request for the additional n1 units in Karp does not affect the high watermark or the low watermark in Karp. As a result, requesting an

additional $n1$ units in Karp cannot read on the requested increase in the minimum resource specification corresponding to the one logical partition recited in claim 10. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

The total allocation of resources in Karp does not read on the sum of all minimum resource specifications. Here again the examiner has failed to identify which teaching in Karp allegedly reads on the minimum resource specification. Recall the minimum resource specification was explicitly defined earlier in claim 10 to specify minimum resources that must be available for a corresponding logical partition to function correctly. Nowhere does the total allocation of resources in Karp read on the sum of minimum resource specification, where each minimum resource specification specifies resources that must be available for a corresponding logical partition to function correctly. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

The high watermark in Karp does not read on the total resources that are permanently enabled in the computer system. As discussed in detail above, Karp is devoid of any teaching regarding total resources that are permanently enabled in the computer system. The high watermark in Karp is used by a threshold to determine whether to allocate $n1$ additional units to a task based on the soft and hard limits that pertain to a user running the task. The high watermark in Karp thus does not read on the total resources that are permanently enabled in the computer system, as recited in claim 10. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

Claim 10 further recites:

... if the requested increase in the minimum resource specification corresponding to the one logical partition will not cause the sum of all minimum resource specifications for all logical partitions to exceed the

total of resources that are permanently enabled in the computer system,
performing the steps of:
 permitting the requested increase in the minimum resource
 specification corresponding to the one logical partition; ...

The “if” clause above is similar to the clause discussed in the preceding paragraphs. Karp does not teach or suggest these limitations for the many reasons discussed above. While Karp does teach permitting a requested increase in resource allocation to a particular user, nowhere does Karp teach or suggest permitting a requested increase in the minimum resource specification that specifies minimum resources that must be available for a corresponding logical partition to function correctly, as expressly recited in claim 10. For this reason alone, claim 10 is allowable over the combination of Karp, Vega and Circenis.

In the rejection of claim 10 under the combination of Karp, Vega and Circenis, the examiner failed to address the limitation “one of the plurality of logical partitions requesting to increase a minimum resource specification corresponding to the one logical partition.” As a result, the examiner has failed to establish a prima facie case for obviousness for claim 10 under 35 U.S.C. §103(a) based on the combination of Karp, Vega and Circenis.

The examiner rejected the majority of limitations in claim 10 as being unpatentable over Karp. In doing so, however, the examiner addresses language not recited in the claims. Throughout the rejection the examiner replaces “logical partition” with “tasks”. Nowhere do the claimed limitations mention tasks as rejected by the examiner. The examiner has failed to address the express limitations of claim 10 and as a result has failed to establish a prima facie case of obviousness under 35 U.S.C. §103(a). The examiner tries to plug the hole in the rejection by stating:

Karp does not specifically teach each task resides in a logical partition...However, Vega teaches logical partitions running applications and tasks for the purpose of accommodating multiple guests on a single computer. It would have been obvious to one having ordinary skill in the

art at the time of the applicant's invention to modify the teachings of Karp with logical partitions running applications and tasks, as taught by Vega, because it allows for the accommodation of multiple guests on a single computer.

Karp teaches allocating portions of a resource between tasks based on defined limits on users. Vega teaches logical partitions running applications and tasks for the purpose of accommodating multiple guests on a single computer. Applying the teachings of Vega to Karp as suggested by the examiner would not have been obvious to one of ordinary skill in the art at the time the invention was made. A simple example will illustrate. Let's assume the logically-partitioned computer system of Vega is combined with Karp. In such a system, let's assume a particular user uses one particular logical partition. Because the soft and hard limits in Karp apply to users, one could apply the teachings of Karp within a single logical partition in Vega. However, it would not have been obvious to one of ordinary skill in the art at the time the invention was made to apply the teachings of Karp across the logical partitions in Vega. This would require some mechanism for tracking users across logical partitions, which is not taught or suggested in either Karp or Vega. Thus, the combination of Karp and Vega as suggested by the examiner would require additional pieces that are not present in either. As a result, the examiner's combination of Karp and Vega is in error.

Claim 10 recites allocating a temporary resource to a logical partition. A task is not equivalent to a logical partition, and substituting one for the other would not have been obvious to one of ordinary skill in the art at the time the invention was made. The examiner's cavalier attempt to equate the task-based functions of Karp to the logical partitions in Vega requires leaps of logic that can only be filled by an examination of appellant's claims, which amounts to impermissible hindsight reconstruction. It would not have been obvious to one of ordinary skill in the art to port the teachings about dividing up resources between tasks as taught by Karp to allocating resources to logical partitions and controlling access to the minimum resource specification in the unique manner expressly recited in claim 10.

For the many reasons given above, claim 10 is allowable over Karp, Vega, and Circenis. Appellant respectfully requests the examiner's rejection of claim 10 under 35 U.S.C. §103(a) based on the combination of Karp, Vega, and Circenis be reversed.

Claim 11

Claim 11 depends on claim 10 which is allowable for the reasons given above. As a result, claim 11 is allowable as depending on an allowable independent claim. Appellant respectfully requests the examiner's rejection of claim 11 under 35 U.S.C. §103(a) based on the combination of Karp, Vega, and Circenis be reversed.

CONCLUSION

Claims 10-11 are addressed in this Appeal. For the numerous reasons articulated above, appellant maintains the rejections of claims 10-11 are erroneous.

Appellant respectfully submits that this Appeal Brief fully responds to, and successfully contravenes, every ground of rejection and respectfully requests that the final rejection be reversed and that all claims in the subject patent application be found allowable.

Respectfully submitted,

By /derekpmartin/
Derck P. Martin
Reg. No. 36,595

MARTIN & ASSOCIATES, L.L.C.
P.O. Box 548
Carthage, MO 64836-0548
(417) 358-4700
Fax (417) 358-5757

CLAIMS APPENDIX

1-9 (Cancelled)

10. A computer-implemented method for providing at least one temporary resource on demand for a specified resource-time in a computer system that includes a plurality of logical partitions, the method comprising the steps of:

- requesting an enablement code corresponding to the at least one temporary resource for the specified resource-time;
- receiving the enablement code;
- enabling the at least one temporary resource for the specified resource-time;
- using the at least one temporary resource for the specified resource-time;
- reading a minimum resource specification for all of the plurality of logical partitions, each minimum resource specification specifying minimum resources that must be available for a corresponding logical partition to function correctly;
- one of the plurality of logical partitions requesting to increase a minimum resource specification corresponding to the one logical partition;
- determining a total of resources that are permanently enabled in the computer system;
- determining whether the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all of the plurality of logical partitions to exceed the total of resources that are permanently enabled in the computer system;
- if the requested increase in the minimum resource specification corresponding to the one logical partition will cause a sum of all minimum resource specifications for all logical partitions to exceed the total of resources that are permanently enabled in the computer system, denying the requested increase in the minimum resource specification corresponding to the one logical partition; and
- if the requested increase in the minimum resource specification corresponding to the one logical partition will not cause the sum of all minimum resource specifications for

all logical partitions to exceed the total of resources that are permanently enabled in the computer system, performing the steps of:

- permitting the requested increase in the minimum resource specification corresponding to the one logical partition; and
- enabling the at least one temporary resource;
- recovering the at least one temporary resource when the specified resource-time expires.

11. The method of claim 10 further comprising the step of evaluating the enablement code to determine whether the code is valid, wherein the enablement code includes the specified resource-time.

12-24. (Cancelled)

EVIDENCE APPENDIX

An Evidence Appendix is not required for this Appeal Brief.

RELATED PROCEEDINGS APPENDIX

A Related Proceedings Appendix is not required for this Appeal Brief.